arXiv:1604.02694v1 [cs.SI] 10 Apr 2016

Predicting Social Status via Social Networks: A Case Study on University, Occupation, and Region

Hao Fu*, Xing Xie†, Yong Rui†, Defu Lian‡, Guangzhong Sun*, Enhong Chen*
*University of Science and Technology of China
†Microsoft Research
‡University of Electronic Science and Technology of China

Abstract

Social status refers to the relative position within the society. It is an important notion in sociology and related research. The problem of measuring social status has been studied for many years. Various indicators are proposed to assess social status of individuals, including educational attainment, occupation, and income/wealth. However, these indicators are sometimes difficult to collect or measure.

We investigate social networks for alternative measures of social status. Online activities expose certain traits of users in the real world. We are interested in how these activities are related to social status, and how social status can be predicted with social network data. To the best of our knowledge, this is the first study on connecting online activities with social status in reality.

In particular, we focus on the network structure of microblogs in this study. A user following another implies some kind of status. We cast the predicted social status of users to the "status" of real-world entities, e.g., universities, occupations, and regions, so that we can compare and validate predicted results with facts in the real world. We propose an efficient algorithm for this task and evaluate it on a dataset consisting of 3.4 million users from Sina Weibo. The result shows that it is possible to predict social status with reasonable accuracy using social network data. We also point out challenges and limitations of this approach, e.g., inconsistence between online popularity and real-world status for certain users. Our findings provide insights on analyzing online social status and future designs of ranking schemes for social networks.

1 Introduction

Social status is the relative position of a person or group within the society. It helps us to understand the stratification of society, and serves as reference for sociology and related research. Social status itself is an abstract concept. For many years, sociologists have discussed how to measure social status in a given society. Various indicators are proposed to assess social status, among which educational attainment, occupation, and income/wealth are the most commonly adopted (Adler et al. 2000). However, these indicators are expensive to collect at a large scale and sometimes difficult to measure. For example, to measure the socio-economic status of occupations, 73,901 respondents were surveyed in (Ganzeboom and Treiman 1996).

We are interested in alternative ways of deriving social status. The growing popularity of online social networks enables the possibility of measuring social status at a large scale. As users interact with each other in social networks, certain traits are exposed through their online activities, including social status. For instance, in a microblogging website, the number of followers is often used as an obvious indicator of status. Users with many followers are considered prestigious and influential (Hopcroft, Lou, and Tang 2011). In previous works, structural stratification of the Twitter network (Hopcroft, Lou, and Tang 2011) and signed networks (Leskovec, Huttenlocher, and Kleinberg 2010) has been observed. However, it is unclear how it is related to social status in reality. In this study, we focus on the impact of network structure on social status and leave other aspects (e.g., content of tweets) for future work.

We investigate the feasibility of predicting social status via online social networks in the following way. We develop an algorithm that is expected to predict social status, and then compare the results with commonly adopted indicators of social status, e.g., educational attainment, occupation, and income/wealth. If a consistent alignment is observed, we may presume that social status can be successfully predicted with social network data.

We consider educational attainment and occupation, but omit income/wealth in this study due to privacy issues in our dataset. Educational attainment can be measured with the highest degree (Adler et al. 2000), years in school, or the prestige of school that one is affiliated to or graduated from (Hevenstone 2008). We use the prestige of university in this study. Various measures of occupational status have been developed in the past decades. We leverage the results of (Ganzeboom and Treiman 1996) to measure occupational status. Additionally, we also consider the region that one lives in as an indicator of social status. As the development of regions within a country is usually unbalanced, it could yield difference of social status of residents.

For prediction, we follow the methodology of (Ganzeboom and Treiman 1996) where scores of respondents are aggregated to status scores of occupations. In our case, we first measure social status of individual users based on network structure, and then cast it to "status" of real-world en-

This work was done when the first author was visiting Microsoft Research.

tities, namely universities, occupations, and regions. An advantage of this approach is that one may easily compare the status scores of the entities with corresponding indicators and directly observe the difference.

We notice that people are the building blocks of the entities in question. A university can be regarded as a group of affiliated students and employees. An occupation can be represented by workers doing the same job. A region consists of residents who live in it. In many websites, users may specify various information about themselves, e.g., universities that they attended, current/former occupations, and regions that they live in. We denote such information as categorical attributes. Universities, occupations, and regions are represented as *groups* of users. *Members* of a group share the same attribute. We refer to such user-group relations as *membership*. The notion of membership is an abstraction of user attributes, and should not be confused with terms in community detection literature. We use the term *group status* to denote the generalization of social status of group members. It capture the similarity of its members in terms of social status. For simplicity, we represent it as a real-valued measure in this study. A high value indicates that the members have high social status. For example, professionals, when viewed as a whole group, are generally considered more prestige than people with other occupations. The measures of occupational status in (Ganzeboom and Treiman 1996) can be considered as an instance of group status.

A fundamental obstacle to this study is the noise and bias in social network data. This is deeply rooted in the nature of online social networks, where users can specify whatever they want about themselves. One may pretend to be a graduate of a famous university for better public recognition, or provide wrong information for privacy concern. Even if one can distinguish some users who honestly specified their profiles, it is still unclear if they represent the average population, or they are just those who try to promote themselves.

To obtain reliable profiles, we only consider *verified* users who are required to provide proof of identity to the website. A drawback is that verified users are mostly influential and popular. Using only verified users would lead to biased results. We handle this issue by proposing a propagation algorithm that infers the membership of remaining users, expecting to enrich the set of reliable profiles. Although users can specify anything in their profiles, they have no direct control on the profiles of friends, which possibly reveal their real attributes (Mislove et al. 2010).

We make two major contributions in this paper.

- We propose a method that predicts group status based on social network structure, consisting of two steps. For the first step, we propose an efficient propagation algorithm to infer missing membership. For the second step, we discuss possible measures of social status of individuals and how group status is calculated.
- We evaluate our approach on a dataset consisting of 3.4 million users and three types of membership (university, occupation, and region). We compare the predicted status with various facts, and find good correlation between social network and social status in reality. However, inconsistence for certain types of users is also observed.

We first give an introduction of our dataset and a concrete definition of our problem in Section 2. The two steps of our method are described in Section 3 and Section 4 respectively. We then evaluate our method and discuss the results in Section 5. Related works are discussed in Section 6. A discussion of the findings is presented in Section 7.

2 Preliminary

2.1 Dataset

We collected our dataset from Sina Weibo, a popular microblogging website in China. We crawled our dataset in May, 2014 using the website's API. We applied the following strategy to obtain a reasonably "good" sample (Leskovec and Faloutsos 2006) from the website. We first made a random sample of tweets posted during April and May, 2014, from the public timeline, expecting to collect a uniform sample of active users. After removing duplicate authors, we collected 49,719 unique users as seeds. We crawled the 2-hop followees of the seed users to obtain a social graph.

We then crawled profiles of the users, among which 8.7% users were verified. 52.6% of the verified users were verified as individuals, while the others were non-individuals, e.g., government, enterprise, and news media. We removed non-individual users since we study social status of real human. Finally, we obtained a graph with 3.4 million nodes and 475 million directed edges.

The website's API provides very limited data of users. It does not contain any information about educational attainment or occupation, though they are accessible via user's profile page. We further crawled the profile pages of verified individual users. Due to the access limit of the website, we were only able to crawl 63,346 valid pages, covering 39% of the verified individual users. We extracted the membership from profiles and pages as follows.

University Users can specify the universities that they attend or graduate from. A university may have multiple synonymous names. For example, users from Peking University may specify "Peking U." or simply "PKU". We manually inspected all distinct university names and replaced them with full names, so that synonymous names were mapped to the same university. We discarded universities having less than 20 members, because inference on such small samples is unreliable. We ended up with 10,256 users who specified their universities. 158 distinct universities were involved, most of which are located in China.

Occupation Users can specify their occupational titles in profiles. An occupation may have multiple synonymous titles, e.g., programmer and software engineer. In addition, the occupational titles are sometimes ambiguous. One may specify "engineer", but it is unclear if it means software engineer, mechanical engineer, or something else. To solve this problem, we referred to the International Standard Classification of Occupations (ISCO) (Ganzeboom and Treiman 1996) to unify the occupational titles. ISCO is a hierarchical classification system mostly derived from skill requirements at the expense of industry distinctions. It contains 9 major groups and three further levels: 28 sub major groups, 116 minor groups, and 290 unit groups.

We inspected the top 1,000 frequent occupational titles in our dataset, and manually mapped them to ISCO groups. Ambiguous occupational titles were difficult to be mapped to fine-grained levels, so we focus on the first level of ISCO (9 groups). The 1,000 occupational titles were classified to 7 major groups. We ended up with 10,582 users that were associated to ISCO groups.

Region Sina Weibo allows users to specify the regions that they live in. The regions are organized according to the administrative divisions of China, including 4 cities, 28 provinces, and 2 special administrative regions. We followed this division and ignored oversea users. We obtained 72,252 known users spreading in the 34 regions.

We are concerned about the sampling bias of our dataset. For each of the three types of membership, over 84% users are associated to only one group. We find that the distribution of users in these groups is unbalanced. Professional occupations account for 32% users. About half users live in two major cities, Beijing and Shanghai. We checked with the website's statistics using its filtered search, and find the distributions are roughly consistent, so we believe our dataset is a representative sample of the website.

2.2 Problem Formulation

We model a microblogging website as a directed graph $G = (V, E)$. Each node in V corresponds to a user in the website. For simplicity, we use the terms *user* and *node* interchangeably. A directed edge (u, v) is included in E if and only if user u follows user v. We refer to the user being followed as a *followee* and the other user as a *follower*. If two users are following each other reciprocally, they are called *friends*. We represent the sets of followees, followers, and friends of u as $N_O(u)$, $N_I(u)$, and $N_R(u)$ respectively.

We denote $H = \{h_1, h_2, ..., h_m\}$ as the set of groups, where m is the number of groups. For example, H may indicate all the distinct universities in the dataset. The membership of a user u is denoted as H_u, which is a subset of H. User u is a member of group h_k, if $h_k \in H_u$. Users with known membership are denoted as V^+, and unknown users are denoted as V^-. This definition captures the case that a user is associated to multiple groups.

We formulate the prediction problem as follows. We first calculate the status score P_u for each $u \in V$. Given membership of partial known users, we then compute the group status score π_i for each $h_i \in H$. P_u and π_i assess the social status of users and groups. High values indicate high status. The group status is expected to align with the indicators of social status in reality.

3 Membership Inference

As the first step, we introduce how we infer unknown membership from verified users. We start by observing homophily in the dataset. Based on the observation, we propose a supervised propagation algorithm. We then discuss how the model is trained efficiently.

3.1 Homophily

In social networks, homophily refers to the property that users with similar characteristics tend to be associated with each other. Previous works have discovered that users are more likely to follow each other reciprocally if they share similar topics (Weng et al. 2010) or friends (Hopcroft, Lou, and Tang 2011). We wonder if members of the same group are likely to follow each other.

Table 1: Probability of being members of the same group. The standard error of estimated probability is less than 10^{-4}.

User pair	University	Occupation	Region
Reciprocal	**0.135**	**0.483**	**0.530**
One-way	0.075	0.468	0.381
Disconnected	0.018	0.287	0.240

Given a pair of users who follow each other reciprocally, we calculate the probability that they are members of the same group. We also calculate this probability when there is only a one-way "follow" between the two users, as well as when they do not follow each other. The result (Table 1) shows that connected users are more probable to be in the same group than those disconnected. This tendency is more obvious if the users follow each other reciprocally. This observation confirms the homophily and reveals the relation between graph structure and membership.

3.2 Propagation

We represent the strength of membership of an unknown user u as $Q_u = (Q_{u1}, Q_{u2}, ..., Q_{um})$. Q_{ui} denotes the probability of u being a member of group h_i. The membership of a known user v is represented as Q^0_v in a similar way that

$$Q^0_{vi} = \begin{cases} 1 & h_i \in H_v \\ 0 & h_i \notin H_v \end{cases}$$

According to homophily, one may infer a user as a member of the most probable group that her friends belong to. By regarding the newly inferred users as known, the membership of more users can be inferred. This suggests a propagation algorithm for membership inference. Various efficient propagation algorithms have already been proposed for similar tasks (Raghavan, Albert, and Kumara 2007; Pennacchiotti and Popescu 2011). However, little attention has been paid to the strength of social ties. One may share a lot in common with a close friend but know nothing about an acquaintance.

In recent years, a family of supervised random walk algorithms have been proposed to improve the quality of link prediction (Backstrom and Leskovec 2011) and web page ranking (Gao et al. 2011). An edge weight function is learned to bias the random walk. We leverage the techniques developed for supervised random walk to learn social tie strength efficiently, aiming to improve the accuracy of membership inference in large-scale networks.

In our scenario, we consider the social tie between two users is strong if they share the same membership. A close friend should have more impact on the inference. We denote the social tie strength between u and v as a non-negative function $f(u, v)$. Note that $f(u, v)$ is not symmetrical, i.e., $f(u, v) \neq f(v, u)$. While u regards v as a good friend, v

may not treat u in the same way. This leads to a weighted propagation:

$$\hat{Q}_u = \frac{\sum_{v \in N_R(u)} f(u,v) Q_v}{\sum_{v \in N_R(u)} f(u,v)} \quad (1)$$

We calculate $\hat{Q}_u$ for all $u \in V^-$, and replace Q_u with the newly calculated $\hat{Q}_u$. This process is repeated until all Q_u converge. Q_u is initially assigned with $1/2$ for $u \in V^-$. For known users V^+, we fix $Q_u = Q_u^0$, so that their membership is propagated to their friends repeatedly.

3.3 The Model

We assume that $f(u,v)$ depends on the features $X_{uv} = (x_1^{uv}, x_2^{uv}, ..., x_k^{uv})$. In our experiments, we used 11 features, which are the number of followers, the number followees, the number of friends, PageRank, and reversed PageRank (reversing the direction of edges) of both users, as well as the number of common friends. We omit the analysis of features due to space limit. We follow the definition in (Backstrom and Leskovec 2011) and define the strength $f(u,v)$ as a function of the features:

$$f(u,v) = \frac{1}{1 + \exp\{-X_{uv} \cdot \mathbf{w}\}} \quad (2)$$

While other forms of $f(u,v)$ are possible, we find this definition works well in our experiments.

We determine the parameter $\mathbf{w} = (w_1, w_2, ..., w_k)$ from the dataset. We randomly split the known users V^+ into two sets, namely *seed nodes* V_S and *target nodes* V_T. We start with some random value of $\mathbf{w}$. We pretend that the membership of target nodes V_T is unknown. Equation (1) is applied iteratively to infer the membership of unknown nodes and target nodes. We then check the estimation error between the inferred membership $\hat{Q}_u$ and the actual membership Q_u^0 of target nodes, and update $\mathbf{w}$ so that the error is reduced. The above process is repeated until the estimation error is minimized. This idea is formulated as an optimization problem that minimizes the loss function

$$L(\mathbf{w}) = \frac{1}{2} \sum_{u \in V_T} \left\| \hat{Q}_u - Q_u^0 \right\|^2 + \frac{\mu}{2} \|\mathbf{w}\|^2 \quad (3)$$

A regulation term is included to avoid overfitting. We leverage the techniques in (Gao et al. 2011) to solve the optimization problem efficiently. Note that the propagation described by Equation (1) is equivalent to

$$\arg\min_{Q} \sum_{u \notin V_T} \left\| \hat{Q}_u - Q_u \right\|^2$$

A new loss function is obtained by including the above term

$$\begin{aligned} L(\mathbf{w}, Q) = & \frac{1}{2} \sum_{u \in V_T} \left\| \hat{Q}_u - Q_u^0 \right\|^2 + \\ & \frac{\lambda}{2} \sum_{u \notin V_T} \left\| \hat{Q}_u - Q_u \right\|^2 + \frac{\mu}{2} \|\mathbf{w}\|^2 \end{aligned} \quad (4)$$

Now that the membership strength Q is treated as an argument of the loss function. It is determined by solving the optimization problem rather than propagating using Equation (1). We apply a stochastic gradient descent algorithm to optimize $L(\mathbf{w}, Q)$. Q and $\mathbf{w}$ are updated simultaneously. With a single server, it takes about 40 iterations in less than 1 hour to converge on our dataset. We then obtain the membership strength Q immediately.

4 Predicting Status

Now we proceed to predict group status based on the inferred membership. We first discuss measures of social status of users. We then discuss how group status is predicted.

4.1 Measuring Status

A theory of status in social networks was implicit in (Guha et al. 2004) and later developed in (Leskovec, Huttenlocher, and Kleinberg 2010). The status theory assumes a status ordering of users and that directed links imply the status ordering. In a microblogging website, a link (u,v) can be explained as u having a lower status than v has. If u regards v as having a higher status, while w is considered more prestige than v, the status theory predicts that w's status is also higher than u's.

We examine if the Weibo network can be explained by the status theory. The status theory makes prediction on the type of triangles. Given that (u,v) and (v,w) both exist, it predicts that the link between u and w must be (u,w) rather than (w,u). We search for evidence by counting the number of triangles. A triangle of type I consists of (u,v), (v,w), and (u,w), which is a positive evidence of the status theory. A triangle of type II is a directed circle consisting of links through u, v, and w, which is a negative evidence. It can be shown that every triangle that contains exactly three links is isomorphic to either of the two triangles. In our dataset, the ratio between the two types is 5.8, which means about five times more triangles of type I is presented in the network than type II. This clearly shows that the Weibo network can be well explained by the status theory.

In (Leskovec, Huttenlocher, and Kleinberg 2010), a heuristic was introduced to order users from high status to low status. It starts with a random ordering of nodes, and swap the positions of randomly selected nodes repeatedly, aiming to increase the number of links from low status to high status. In a graph with millions of nodes, we find it rather time-consuming to find a good solution, so we consider other measures that can be efficiently computed, including the number of followers, eigenvector centrality, and PageRank. It is worthwhile to investigate and design more complicated measures, but it is beyond the scope of this paper. These measures follow the similar intuition that a user is more prestige than her followers. Note that these network-based measures are not measures of social status yet, until we validate with indicators of social status in reality.

The number of followers is a natural and obvious measure of a user's position in the social network. A user with millions of followers has direct influence on that many users, so one may assume that she has a high social status. However, this measure does not accurately capture the status, e.g., it can be easily manipulated (Ghosh et al. 2012b).

Eigenvector centrality (Bonacich 1987) measures the importance of a user in a social network. It improves the number of followers by considering the structure of the whole network. It is defined based on the concept that a user is as important as those who she has influence on. Specifically, the eigenvector centrality of a user is proportional to the sum of eigenvector centrality of her followers. Let A be the adjacency matrix of the graph $G = (V, E)$, i.e., $A_{uv} = 1$ if $(u, v) \in E$, or $A_{uv} = 0$ otherwise. Eigenvector centrality is essentially the principal eigenvector of A^T. Given that a normal user is only capable to read a limited number of tweets per day, a user having too many followees is unlikely to be influenced. Such users do not contribute much to the influence of their followees. A limitation of eigenvector centrality is that it does not consider such cases.

PageRank has been widely used to rank web pages. It is adopted for ranking users in social networks (Tang et al. 2009; Weng et al. 2010). It improves eigenvector centrality by weighing nodes with degrees. Users following many others contribute less to their followees. PageRank is defined as the stationary distribution of a random walk process:

$$P_v = d \cdot \sum_{u \in N_I(v)} \frac{P_u}{|N_O(u)|} + \frac{1-d}{|V|} \tag{5}$$

The damping factor d is usually assigned with 0.85. PageRank is shown to be a good measure of status (Hopcroft, Lou, and Tang 2011) and influence (Weng et al. 2010) in microblogs, so we focus on it in this paper.

4.2 Group Status

Now we discuss how status score P of users is transformed to group status score π. A natural assumption is that a group's status is determined only by its members. Two groups may interact with each other via the links between their members, but such interactions were supposed to be captured by P. Recall that we represent the membership of users as a set of probabilities Q. If we know u as a member of h_i for sure, i.e., $Q_{ui} = 1$, π_i should be affected by P_u to the maximum extent. If we are uncertain, e.g., $Q_{ui} = 1/2$, less impact to π_i should be expected.

Another assumption is that a group's status should be irrelevant to its size. First, the number of members is not a dominating factor of the status of a group. For example, most university rankings focus on the quality and reputation of higher education, rather than the number of students and employees. Second, membership of only a sample of users is known in this study. Even with the membership inference, we do not know for sure the exact number of actual members of a group. Third, the number of active social network users may vary among groups. Groups with more active users may get higher status scores. By canceling the number of social network users, we may partially discard such bias.

Putting the above assumptions together, we define a group's status score as the ratio between the expected sum of its members' status score and its expected size, i.e.,

$$\pi_i = \frac{\sum_{v \in V} Q_{vi} P_v}{\sum_{v \in V} Q_{vi}} \tag{6}$$

Table 2: Accuracy of membership inference

	Method	Accuracy	Avg. Accuracy
University	SP	0.245	**0.199**
	UP	**0.246**	0.135
	LP	0.086	0.015
	Hybrid	0.124	0.094
Occupation	SP	**0.399**	**0.287**
	UP	**0.399**	0.210
	LP	0.392	0.187
	Hybrid	0.385	0.252
Region	SP	**0.760**	**0.522**
	UP	0.643	0.321
	LP	0.531	0.151
	Hybrid	0.665	0.477

The contribution to π_i of a user u is proportional to u's membership strength. A ranking can be generated by ordering groups in descending order of π_i.

5 Experiments

In this section, we evaluate the performance of the proposed membership inference algorithm. And then we compare and discuss the predicted group status with reality.

5.1 Membership Inference

Settings Membership inference estimates unknown membership from known users. Given a type of membership, we split the known users into two sets, namely training set and testing set. We feed our algorithm with the training set, and evaluate the accuracy of inferred membership on the testing set. We consider the most probable group ($\arg\max_{h_i \in H} Q_{ui}$) as the inferred group. The inference is considered correct if the user is an actual member of the inferred group. We define the accuracy as the fraction of correctly inferred users. As the group size is unbalanced, we also consider a balanced version of accuracy, which is the average over the accuracy on each group.

Our approach, supervised propagation (SP), requires seed nodes and target nodes. We keep 80% of the training set for seed nodes, and leave the other for target nodes. The parameter λ trades off between the inference error on target nodes and the convergence of propagation. The parameter μ controls the extent of overfitting. We find that setting $\lambda = 1$ and $\mu = 1$ works well in most cases. We repeat the iteration until the relative improvement of $L(\mathbf{w}, Q)$ is less than 1%. We also evaluated the following approaches for comparison.

Uniform Propagation (UP) To investigate the effect of the learned social tie strength $f(u, v)$, we consider a simplified model where $f(u, v)$ is fixed to a constant for each $(u, v) \in E$, so that the membership is propagated uniformly.

Label Propagation (LP) A clustering algorithm described in (Raghavan, Albert, and Kumara 2007) starts with a set of known nodes and update the label of other nodes iteratively. Every node takes the most frequent label in its neighbors. This algorithm can be considered as a "hard" version of UP, as membership is indicated with categorical labels rather than probabilities.

Table 3: AUC of predicting top universities

Method	Project 211	Project 985	C9 League
SP	**0.612**	**0.770**	**0.903**
UP	0.601	0.767	0.882
PR	0.431	0.618	0.776

Hybrid Pennacchiotti and Popescu (2011) proposed a hybrid algorithm that infers binary labels on Twitter. They first apply a machine learning algorithm to infer node labels individually. The inferred label is then propagated via edges to correct previous errors. The original algorithm can only handle binary attributes. We decompose the membership of m groups into m binary attributes, and apply the algorithm separately. To make a fair comparison, only profile features and social network features are used.

We also considered other recent works (Backstrom, Sun, and Marlow 2010; Mislove et al. 2010; Zhao et al. 2013; Li, Wang, and Chang 2014) that might be used for membership inference. A discussion of them is presented in Section 6.

Results We use a 10-fold cross validation in the experiments. The results are reported as the averages (Table 2). UP performs better than LP, indicating the advantage of representing membership as probabilities instead of discrete labels. Hybrid performs quite well on inferring occupations and regions, but not as good as SP. An explanation is that structural information is not fully utilized in Hybrid, since it propagates labels only to 1-hop neighbors. SP achieves the highest accuracy on the three types of membership. This shows the power of the learned social tie strength $f(u, v)$. We omit detailed discussion of the results here, since it is not the main focus of this paper.

5.2 Prediction

Now we have the inferred membership strength Q. We first calculate status scores P of users, and then put them together to obtain group status scores π as described in Section 4. Membership inference is expected recover missing membership so that we can have better prediction of group status. Our first concern is about its ability to improve the quality of prediction. Our second concern is that whether the predicted group status is consistent with social status in the real world. If not, when and why?

For the first concern, we evaluate approaches with different settings of membership inference. We compare the results with facts collected from reality across domains. For the second concern, we distinguish cases when inconsistence is observed and conduct analysis on possible factors. The following approaches are evaluated.

SP This is the major approach proposed in this paper, where missing membership is inferred with the supervised propagation algorithm.

UP We consider a variant of our method in which uniform propagation is used for membership inference. We compare with this approach to validate that improving membership inference has a positive effect on predicting group status.

Table 4: Top 20 universities by group status

#	SP	PR
1	[123]Peking Univ.	Cheung Kong Graduate School of Business
2	[1*]Communication Univ. of China	Southwest Univ. of Political Sci. and Law
3	[*]Central Academy of Drama	[*]Shanghai Theatre Academy
4	[123]Tsinghua Univ.	[1]China Univ. of Political Sci. and Law
5	[*]Shanghai Theatre Academy	[*]P.L.A. Arts College
6	[*]Beijing Film Academy	[*]Central Academy of Drama
7	[12]Renmin Univ. of China	Univ. of International Relations
8	Tianjin Normal Univ.	[1]Hunan Normal Univ.
9	[123]Fudan Univ.	[*]Sichuan Fine Arts Institute
10	[1]China Univ. of Political Sci. and Law	China Europe Int'l Business School
11	[12]Beijing Normal Univ.	[123]Peking Univ.
12	[12]Huazhong Univ. of Sci. and Tech.	[1]Univ. of Sci. and Tech. Beijing
13	[123]Nanjing Univ.	[*]Beijing Film Academy
14	[12]East China Normal Univ.	Tianjin Normal Univ.
15	[123]Shanghai Jiao Tong Univ.	Guangdong Univ. of Tech.
16	[123]Zhejiang Univ.	[1*]Central Conservatory of Music
17	[12]Nankai Univ.	[123]Fudan Univ.
18	China Europe Int'l Business School	Capital Normal Univ.
19	[123]Univ. of Sci. and Tech. of China	Chinese Academy of Social Sciences
20	Southwest Univ. of Political Sci. and Law	[12]Nankai Univ.

[1]Project 211, [2]Project 985, [3]C9 League, [*]Art/music/acting school

PR In this approach, we do not perform membership inference. We calculate PageRank on the whole network and then take a group's status score as the average of its known (verified) members' PageRank. We evaluate this approach to verify the necessity of membership inference.

Evaluation is performed with membership of university, occupation, and region. We discuss the results separately.

University To assess the quality of predicted university status, we compare it with public rankings by third-parties. Multiple Chinese university rankings are published by different institutes, e.g., the ARWU rankings and the QS World University Rankings. These rankings focus on different factors of quality and reputation of university. However, their source data and criteria are rarely published, rendering them less convincing. Therefore, we refer to some kind of "official" classification of top universities. We consider Project 211, Project 985, and the C9 League. Project 211 and Project 985 are both Chinese government-run education projects[1]. Universities designated as Project 211 or Project 985 institutes are supported with first priority. Until 2015, 116 universities and 39 universities are sponsored by the two projects respectively. The C9 League consists of the first nine universities selected for Project 985. It is often referred to as the Chinese equivalent of the US Ivy League. These universities are generally regarded as top universities in China.

We are aware that university ranking is a sensitive topic. A university can argue against a ranking because they are ranked lower than they think. To avoid this, we formulate a binary classification problem instead. Given a university, the task is to predict if it is one of the top universities (Project

[1]http://www.moe.gov.cn/

Table 5: Ranking of occupations by SIOPS, ISEI, and group status

#	SIOPS	ISEI	SP/UP	PR
1	Professionals	Professionals	Legislators, Senior Officials & Managers	Legislators, Senior Officials & Managers
2	Legislators, Senior Officials & Managers	Legislators, Senior Officials & Managers	Professionals	Professionals
3	Technicians & Associate Professionals	Technicians & Associate Professionals	Technicians & Associate Professionals	Technicians & Associate Professionals
4	Craft & Related Trades Workers	Clerks	Service, Shop & Market Sales Workers	Skilled Agricultural & Fishery Workers
5	Clerks	Service, Shop & Market Sales Workers	Clerks	Service, Shop & Market Sales Workers
6	Skilled Agricultural & Fishery Workers	Craft & Related Trades Workers	Craft & Related Trades Workers	Clerks
7	Service, Shop & Market Sales Workers	Skilled Agricultural & Fishery Workers	Skilled Agricultural & Fishery Workers	Craft & Related Trades Workers

211, Project 985, or the C9 League). We do this by sorting universities in descending order of group status score.

We adopt the standard notion of true positive rate and false positive rate to assess the performance. The true positive rate (TPR) is defined as the fraction of correctly identified top universities out of the actual top universities. The false positive rate (FPR) is defined as the fraction of misclassified non-top universities out of the actual non-top universities. The trade-off between TPR and FPR is visualized by the ROC curve. The overall performance is quantified by the area under the curve (AUC).

The performance of prediction is presented in Table 3. C9 League universities can be better predicted than Project 211 and Project 985 universities by the three approaches. In general, the results show positive correlation between the predicted group status and their real status.

The PR approach, which simply takes the average of known users' PageRank without membership inference, performs poorly. It is even worse than random ($\mathrm{AUC} < 0.5$) for Project 211 universities. We find that users from non-top universities are less likely to have their accounts verified. For the three classifications, the average number of verified users from non-top universities is less than 58, while the similar average among top universities is greater than 84. To understand this behavior, we manually inspect user profiles from top-ranked universities by PR. We find a tendency that graduates of top universities prefer to have their accounts verified for better public recognition, while only elite graduates of non-top universities would verify their accounts. In other words, the PR approach compares average graduates of top universities with elite graduates of non-top universities, which does not make much sense.

This confirms our concern about the noise and bias in online data. Although we can obtain reliable profiles by considering only verified users, they turn out to be a biased sample of the total population. Certain methods are required to filter and clean the data. In our case, we propose a propagation algorithm to infer missing membership. The simplified version of the algorithm (UP) improves the performance of prediction significantly, implying that membership inference is an effective way to recover missing data. By learning the social tie strength (SP), the performance is further improved. Although the algorithm can not infer missing membership perfectly (Table 2), it turns out to be sufficient to predict group status reasonably well.

We present the top 20 universities by SP and PR in Table 4. A significant number of art, music, or acting schools are in the top 20. First, artists, musicians, and actors are actually more famous to the general public than other occupations, e.g., most people know Lady Gaga but rarely hear of Paul Erdős. As most of them graduate from art, music, or acting schools, the status of these schools is promoted on social networks. Second, some law schools and business schools are inferred with high status. This can be explained by the occupational status of their graduates, e.g., lawyers and managers. The remaining universities by SP are mostly associated to Project 211, Project 985, or C9 League, while only 7 universities by PR are associated to Project 211.

In summary, we observe sampling bias of verified users, which leads to undesirable prediction of simple approaches (e.g., PR). By inferring missing membership, this issue can be partially corrected. Our result shows that it is possible to prediction status of universities with reasonable accuracy. However, inconsistence is observed for certain types of institutions. Their graduates gain extraordinary occupational influence on social networks by nature, and it is a stubborn problem on predicting social status via social networks.

Occupation In sociology, three indicators are commonly used for occupational status, namely occupational prestige, socio-economic index, and Erikson-Goldthorpe-Portocarero (EGP) class categories. Occupational prestige is a scalar measure of the respect with which an occupation is regarded by others. Socio-economic index scores are calculated as weighted averages of standardized measures of income and educational attainment. EGP class categories are nominal typologies that combine occupational information.

Ganzeboom and Treiman (Ganzeboom and Treiman 1996) developed internationally comparable measures of the above three indicators, namely Standard International Occupational Prestige Scale (SIOPS), International Socio-Economic Index of Occupational Status (ISEI), and EGP class categories. As EGP class categories are not ordered, we only consider the other two here. We compare the above indicators with predicted status.

We order the 7 occupation groups from high status to low status in Table 5. The orderings by SP and UP are the same. In general, ISEI is more correlated to predicted group status than SIOPS, indicating that social status on Sina Weibo appears more relevant to income and educational attainment. Further investigation is required to justify this observation.

SP/UP and ISEI are similar, except that two pairs of adjacent groups are swapped. First, "Professionals" and "Legislators, Senior Officials & Managers" are swapped. We inspected the occupational titles of the two groups. Scientists

Table 6: Spearman correlation between rankings of regions by group status and social-economic statistics

Method	Per capita income	Per capita GDP
SP	0.683	0.498
UP	**0.690**	**0.502**
PR	0.467	0.397

and actors are ranked among the top, but they consist of only 18% professionals. The other group contains only corporate managers. We did not find any legislator or official in the top 1,000 frequent occupational titles. This is possibly due to the sensitivity of their identities. Corporate managers appear to be more popular than other professionals, e.g., journalists, doctors, and engineers. Second, "Clerks" and "Service, Shop & Market Sales Workers" are swapped. Marketing related occupations, e.g., advertising, salesman, and public relation, consist of 85% of the second group. We find evidence of link farming aiming to increase their influence. One method is following back anyone who follows them. Ghosh et. al (Ghosh et al. 2012b) find it an effective way to promote the ranking on Twitter. Marketers follow back 37% of their followers on average, while the similar ratio is 24% and 31% for "Clerks" and all users.

The above observations suggest that the slight difference between SP/UP and ISEI is mainly caused by the unbalanced distribution of known occupations as well as link farming by marketers. The first issue could be solved by considering fine-grained occupational groups, if we have more specified occupational titles in user profiles. For the second issue, anti-spam techniques could be helpful to discard the effect of link farming (Ghosh et al. 2012b). The PR approach is less relevant to ISEI, suggesting the effect of membership inference.

Region Recall that we predict group status according to status of individuals. Correspondingly, we focus on per capita socio-economic statistics, namely per capita income and per capita GDP in 2014[2], which are indicators of average income and wealth of citizens in each region.

We rank the regions according to the statistics and predicted group status. We measure the correlation with Spearman correlation coefficient. A perfect Spearman correlation of +1 or -1 occurs when the two rankings are the same with or the reverse to each other, and 0 indicates irrelevance. The result (Table 6) shows that per capita income is more relevant to predicted group status than per capita GDP. This implies the different nature of the two statistics. Per capita income is a direct measure of people's income, while per capita GDP involves other social and economic factors, which can not be captured by social network data.

SP and UP produce similar rankings, while UP performs slightly better. PR is significantly less correlated to the statistics than SP and UP. To see the difference, we plot the regions in Figure 1. Figures 1(c) and 1(d) show that regions along the coast in Eastern and Southern China are more developed, and Western China is less developed. The result of SP/UP (Figure 1(a)) agrees with this tendency.

[2]http://www.stats.gov.cn/

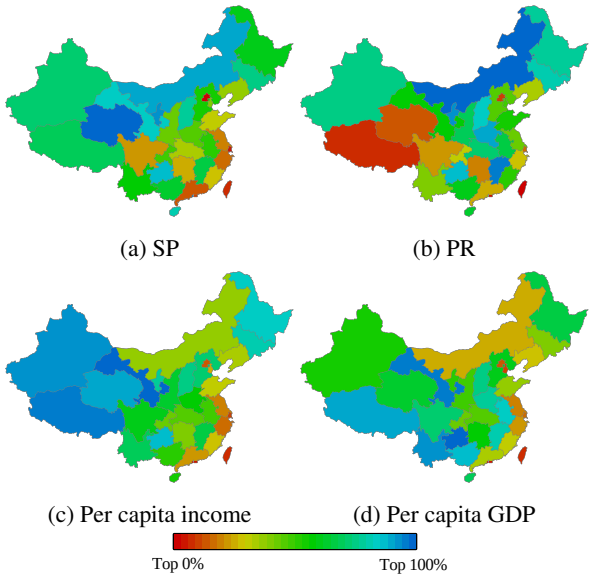

Figure 1: Ranking of regions by group status and socio-economic statistics. UP is similar to SP and thus omitted.

The PR approach assigns Tibet and Qinghai, which are located in Western China, with unusually high status. In our dataset, 67 and 54 verified users are located in Tibet and Qinghai respectively, which are much less than the average number (2,125). We find 13 out of the 121 users are verified as Tibetan Buddhist leaders. We do not observe such a large fraction of religious leaders in other regions. This can be explained by the fact that Tibetan Buddhism spreads mainly in these two regions nowadays. 7 of the Tibetan Buddhist leaders are among the top 20 users in these two regions. When we discard them in the PR approach, Tibet and Qinghai are ranked at #17 and #33 respectively. This shows again the sampling bias in social network data, where verified users may not represent the average population. By inferring unverified users, the bias can be partially fixed.

The top 5 regions in terms of per capita income are Beijing, Shanghai, Hong Kong, Macao, and Taiwan. Four of them are also ranked in the top 5 by SP and UP, but Macao is assigned with surprisingly low status. We find that only 40 verified users are located in Macao. Inference on such a small sample seems unreliable, even if membership inference is applied. A better coverage of users may solve this issue.

6 Related Work

Social status is one of the fundamental principles that describe the stratification of society. A related concept, namely social capital (Putnam 1993), regards one's social ties as a resource that can be utilized to access more resources (Ellison, Steinfield, and Lampe 2007). It is different from social status as the later focuses on the position in society and explains social ties as a consequence of status. A few works have studied social status in social networks. Leskovec et. al (2010) find that the structure of signed networks can be

explained by status theory. Hopcroft et. al (2011) observe that users with similar social status tend to follow each other reciprocally on Twitter. Many works focus on ranking with different principles, e.g., influence. However, little is known about how online status is related to reality. In this work, we try to answer this question with the proposed method and evaluation on three different domains. Our method is built on the ideas and techniques developed for social network influence analysis and user profiling, so we present a brief discussion of them here. We also distinguish our work with a similar problem, namely ranking in heterogeneous network.

Influence in microblogs A number of works focus on analyzing influence in microblogs based on the network structure (Kwak et al. 2010) and content of tweets (Cha et al. 2010). Different models are proposed to measure the influence (Bakshy et al. 2011; Romero et al. 2011). Another body of research focuses on topic-specific influence in microblogs. Given certain topics, the task is to find the most influential users in consideration of relevance and importance (Pal and Counts 2011; Ghosh et al. 2012a; Bi et al. 2014).

Some works start with PageRank to design better ranking models (Tang et al. 2009; Weng et al. 2010; Ghosh et al. 2012b). In our work, we also find it a good measure of social status. We are aware that better models for social status can be developed, and leave it for future work.

Ranking in heterogeneous network A social network with membership information can be modeled as a heterogeneous network, containing two types of nodes: users and groups. A user and a group are linked by an edge if the user is a member of the group. Predicting group status is casted to the ranking problem in the heterogeneous network. Sun et al. (2009) proposed a family of algorithms that perform clustering and ranking simultaneously. They applied the algorithms on DBLP to cluster and rank academic papers and conferences. Ji et al. (2011) developed methods to rank and classify nodes simultaneously.

The above works assume that links in the network are fully presented. In our scenario, only a small fraction of biased links between groups and users are known, so these algorithms can not work properly. In addition, online social networks, where users can establish links casually, are much more noisy than citation networks or co-author networks.

User profiling Membership inference can be considered as a user profiling problem, where missing attributes are inferred from known attributes and network structure. Several works are built on the idea of propagation. Backstrom et al. (2010) proposed an algorithm specialized to locations. Mislove et al. (2010) infer user profiles in a university by applying a graph clustering algorithm. However, members of the detected cluster do not necessarily share the same attribute. Pennacchiotti and Popescu (2011) proposed a hybrid algorithm to infer binary attributes, by combining machine learning methods and propagation on networks. More complicated models are proposed to capture the correlation between different types of attributes (Zhao et al. 2013), and the relationship between node and edge (Li, Wang, and Chang 2014). Limitations of the user profiling problem were discussed in (Cohen and Ruths 2013).

Given that we are processing a network with millions of nodes and hundreds of groups, we need an inference algorithm that is scalable to both the size of network and the number of groups. We notice that little attention has been paid to social tie strength in user profiling. We leverage the techniques developed for supervised random walk (Backstrom and Leskovec 2011; Gao et al. 2011) to learn the strength. It turns to be reasonably efficient and effective.

7 Conclusion

In this paper, we investigate the relationship between social network and social status in reality. Our method consists of two steps, namely membership inference and predicting group status. We compare the predicted status with various indicators from the real world, covering domains of university, occupation, and region. Our results show that it is possible to predict social status via social networks with reasonable accuracy. However, inconsistence between the status in social networks and reality is observed, implying challenges of this problem.

The inconsistence comes in two folds. First, as we use the membership information as a bridge to link social network and the real world, it is unclear what kind of users would provide such information. We find bias in user samples in all three domains. Elite graduates from non-top universities rather than average graduates would have their accounts verified. While officials try to hide their identities, marketers are eager to promote themselves. Religious leaders, who are very prestigious, concentrate in certain regions. Although users can specify whatever profiles they want, they have no direct control on their friends' profiles (Mislove et al. 2010). This suggests membership inference as a potential way to fix such bias. We believe that this should not be a big problem if more comprehensive data is available.

Second, users do not behave equally active in social networks. Graduates from art, music, or acting schools are quite popular on social networks. It is due to the occupational nature of the schools. We also find evidence that marketers try to gain influence by link farming, rendering their online status higher than they have in reality. Our current measure of social status does not take these issues into consideration. We consider this as a challenge for future studies, since it is unclear yet if these behaviors are related and if it is possible to describe them in a unified model.

We believe the study of social status is valuable to several applications beyond the discovery itself. We show that social networks imply social status in reality. Better recommendation systems could be designed for online shopping by considering social status, since it is closely related to income and educational attainment. Social status can also be utilized to analyze trust and influence on social networks. Other types of data could also be investigated, e.g., retweeting/sharing activities, textual contents, and images. The proposed approach may also serve as a tool to analyze social status of other entities, e.g., companies, websites, books, and movies.

References

[Adler et al. 2000] Adler, N. E.; Epel, E. S.; Castellazzo, G.; and Ickovics, J. R. 2000. Relationship of subjective and objective social status with psychological and physiological functioning: Preliminary data in healthy, white women. *Health psychology*.

[Backstrom and Leskovec 2011] Backstrom, L., and Leskovec, J. 2011. Supervised random walks: Predicting and recommending links in social networks. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*.

[Backstrom, Sun, and Marlow 2010] Backstrom, L.; Sun, E.; and Marlow, C. 2010. Find me if you can: Improving geographical prediction with social and spatial proximity. In *Proceedings of the 19th International Conference on World Wide Web*.

[Bakshy et al. 2011] Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone's an influencer: Quantifying influence on twitter. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*.

[Bi et al. 2014] Bi, B.; Tian, Y.; Sismanis, Y.; Balmin, A.; and Cho, J. 2014. Scalable topic-specific influence analysis on microblogs. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*.

[Bonacich 1987] Bonacich, P. 1987. Power and centrality: A family of measures. *American Journal of Sociology*.

[Cha et al. 2010] Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, P. K. 2010. Measuring user influence in twitter: The million follower fallacy. In *Proceedings of the 4th International Conference on Weblogs and Social Media*.

[Cohen and Ruths 2013] Cohen, R., and Ruths, D. 2013. Classifying political orientation on Twitter: It's not easy! In *Proceedings of the 7th International Conference on Weblogs and Social Media*.

[Ellison, Steinfield, and Lampe 2007] Ellison, N. B.; Steinfield, C.; and Lampe, C. 2007. The benefits of facebook "friends:" social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication*.

[Ganzeboom and Treiman 1996] Ganzeboom, H. B., and Treiman, D. J. 1996. Internationally comparable measures of occupational status for the 1988 international standard classification of occupations. *Social science research*.

[Gao et al. 2011] Gao, B.; Liu, T.-Y.; Wei, W.; Wang, T.; and Li, H. 2011. Semi-supervised ranking on very large graphs with rich metadata. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining*.

[Ghosh et al. 2012a] Ghosh, S.; Sharma, N.; Benevenuto, F.; Ganguly, N.; and Gummadi, K. 2012a. Cognos: Crowdsourcing search for topic experts in microblogs. In *Proceedings of the 35th International ACM Conference on Research and Development in Information Retrieval*.

[Ghosh et al. 2012b] Ghosh, S.; Viswanath, B.; Kooti, F.; Sharma, N. K.; Korlam, G.; Benevenuto, F.; Ganguly, N.; and Gummadi, K. P. 2012b. Understanding and combating link farming in the Twitter social network. In *Proceedings of the 21st International Conference on World Wide Web*.

[Guha et al. 2004] Guha, R.; Kumar, R.; Raghavan, P.; and Tomkins, A. 2004. Propagation of trust and distrust. In *Proceedings of the 13th International Conference on World Wide Web*.

[Hevenstone 2008] Hevenstone, D. 2008. Academic employment networks and departmental prestige. *Why Context Matters*.

[Hopcroft, Lou, and Tang 2011] Hopcroft, J.; Lou, T.; and Tang, J. 2011. Who will follow you back?: Reciprocal relationship prediction. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*.

[Ji, Han, and Danilevsky 2011] Ji, M.; Han, J.; and Danilevsky, M. 2011. Ranking-based classification of heterogeneous information networks. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining*.

[Kwak et al. 2010] Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*.

[Leskovec and Faloutsos 2006] Leskovec, J., and Faloutsos, C. 2006. Sampling from large graphs. In *Proceedings of the 12th ACM International Conference on Knowledge Discovery and Data Mining*.

[Leskovec, Huttenlocher, and Kleinberg 2010] Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010. Predicting positive and negative links in online social networks. In *Proceedings of the 19th International Conference on World Wide Web*.

[Li, Wang, and Chang 2014] Li, R.; Wang, C.; and Chang, K. C.-C. 2014. User profiling in an ego network: Co-profiling attributes and relationships. In *Proceedings of the 23rd International Conference on World Wide Web*.

[Mislove et al. 2010] Mislove, A.; Viswanath, B.; Gummadi, K. P.; and Druschel, P. 2010. You are who you know: Inferring user profiles in online social networks. In *Proceedings of the 3rd ACM International Conference on Web Search and Data Mining*.

[Pal and Counts 2011] Pal, A., and Counts, S. 2011. Identifying topical authorities in microblogs. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*.

[Pennacchiotti and Popescu 2011] Pennacchiotti, M., and Popescu, A.-M. 2011. Democrats, republicans and starbucks afficionados: User classification in Twitter. In *Proceedings of the 17th ACM International Conference on Knowledge Discovery and Data Mining*.

[Putnam 1993] Putnam, R. D. 1993. The prosperous community: social capital and public life. *The american prospect*.

[Raghavan, Albert, and Kumara 2007] Raghavan, U. N.; Albert, R.; and Kumara, S. 2007. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* 76.

[Romero et al. 2011] Romero, D. M.; Galuba, W.; Asur, S.; and Huberman, B. A. 2011. Influence and passivity in social media. In *Proceedings of the 20th International Conference on World Wide Web Companion*.

[Sun, Yu, and Han 2009] Sun, Y.; Yu, Y.; and Han, J. 2009. Ranking-based clustering of heterogeneous information networks with star network schema. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining*.

[Tang et al. 2009] Tang, J.; Sun, J.; Wang, C.; and Yang, Z. 2009. Social influence analysis in large-scale networks. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining*.

[Weng et al. 2010] Weng, J.; Lim, E.-P.; Jiang, J.; and He, Q. 2010. TwitterRank: Finding topic-sensitive influential Twitterers. In *Proceedings of the 3rd ACM International Conference on Web Search and Data Mining*.

[Zhao et al. 2013] Zhao, Y.; Wang, G.; Yu, P. S.; Liu, S.; and Zhang, S. 2013. Inferring social roles and statuses in social networks. In *Proceedings of the 19th ACM International Conference on Knowledge Discovery and Data Mining*.